Being cool gets slightly easier

An efficient and exact noncommutative quantum Gibbs sampler
Authors: Chi-Fang Chen, Michael J. Kastoryano, András Gilyén
arXiv:2311.09207

Recommended with a Commentary by Joel E. Moore, University of California, Berkeley, and Lawrence Berkeley National Laboratory

Many of the key discoveries in quantum condensed matter physics were enabled by the ability to cool matter to increasingly small fractions of a Kelvin. Cooling bulk matter to millikelvin temperatures requires a considerable degree of patience–even if a low-temperature reservoir or "bath" is available and in contact with the boundaries of the system, the decrease of thermal conductivity as $T \to 0$ means that a long time is required to get heat out. Emerging quantum computers and emulators are confronting the same problem, and we will review their technical challenges in a bit more detail below. Briefly, the most interesting states in realistic models of materials are largely beyond current quantum computers because the computation cannot reach sufficiently low energy scales in the available coherence time. New algorithms, however, such as the one in the highlighted paper, offer the promise for more rapid "algorithmic cooling," while also connecting to some basic questions about the behavior of quantum matter.

It should be noted that cooling to some ground states is trivial, while cooling to others is likely to remain impossible even on quantum computers. Let us focus on spin models with a spin-half degree of freedom at each site of some lattice. The state with all spins up, which is a ground state of an Ising ferromagnet in zero field, is easy to prepare. Conversely, if random magnetic fields are added to an Ising model, finding ground states is generally NP-complete, which implies that if a polynomial-time algorithm could be found to solve this random-field Ising model problem, an enormous number of other problems would also be solvable in polynomial time [1]. There are problems, such as factoring, for which there are polynomial-time (in the number of digits) algorithms on quantum computers but which are believed to require exponential time on classical computers. No polynomial-time quantum algorithm has been found for an NP-complete problem, however, which is consistent with the widely held belief that quantum computers cannot solve NP-complete problems in polynomial time.* The good news is that reasonable models for an enormous range of materials are neither shown to be impossible nor yet solved in polynomial time, leaving much work to do.

*Finding the ground state of more complicated quantum Hamiltonians is at least as hard: QMA-complete, where QMA, for Quantum Merlin Arthur, is another complexity class. An open question is whether the number of complexity classes defined by computer scientists grows polynomially or exponentially in time.

Quantum computers are already very good at unitary evolution: given an initial state, the time evolution of that state under fairly complicated Hamiltonians can be predicted to high accuracy, even for time-dependent Hamiltonians. This led to the first approach to obtain ground states: prepare the ground state of a simple Hamiltonian, then evolve the Hamiltonian slowly into a more complicated Hamiltonian. The adiabatic theorem of quantum mechanics says that if the evolution is "slow enough," then the system will remain in its ground state. How can this be consistent with the putative exponential difficulty of NP-complete problem such as the random-field Ising model mentioned above? There must be an exponentially small gap along generic paths (at least, along any path that can be reliably found in polynomial time), which would require exponentially long time to pass through adiabatically.

If we were forced to ramp the Hamiltonian over some finite time in a large enough system, we would expect transitions out of the ground state. The degree to which the system departs from the ground state is universal (i.e., dependent only on the universality class) at a quantum critical point: this quantum Kibble-Zurek mechanism [2, 3] is used in practice to measure some critical exponents in quantum simulators. If we wait long enough after the Hamiltonian of the system has reached its final value, we might expect a generic system to locally resemble the Gibbs ensemble $\rho_G = \exp -H/(k_B T)$ at some low temperature T determined by the energy density above the ground state. But again this time for low T is expected to be quite long. So we might ask whether it is possible to reach the low-temperature Gibbs ensemble more rapidly, for example by emulating the refrigeration process used to cool actual materials.

The reader may wonder in passing what it means to call the temperatures of interest in correlated materials "low." We define temperature in a simulation of some effective Hamiltonian H in comparison to the energy scale of the terms in the Hamiltonian. For example, the Hubbard model has hopping term t and interaction term U that are typically both of order electron volts, or tens of thousands of Kelvin. Therefore even the high-temperature superconductors show the superconducting state at quite a low energy scale relative to the starting point, which therefore requires a long computation time to access. The success of quantum computations of matter to date has chiefly been on niche cases where the dynamics is interesting even on the naïve scale of the terms in the Hamiltonian, as in simulations of the toric code Hamiltonian [4] or the Heisenberg spin chain [5].

Imitating the way refrigeration works, we could consider the system of interest as a subsystem of a larger system. Then unitary evolution of the whole system can appear non-unitary and non-energy-conserving from the point of view of the subsystem, which is now "open." We can hope to cool the subsystem to a low-temperature Gibbs ensemble by finding an environment that generates non-Hermitian dynamics of the subsystem. As in conventional cooling, we would like the subsystem dynamics to satisfy a detailed balance constraint. The point of the work of Chen, Kastoryano, and Gilyén (CKG)[6] is to construct such a cooling approach, which can be viewed as an analogue for quantum computers of classical Monte Carlo algorithms, which likewise describe open system dynamics satisfying detailed balance.

The open-system evolution of the subsystem in the CKG approach is fairly technical but has the following important properties. The Gibbs state with density matrix ρ_G is an exact stationary state of the evolution, because it satisfies a detailed-balance constraint. The Lindbladian evolution is relatively efficient to simulate in that the total time taken

scales as the product of $\beta = 1/(k_B T)$ and the mixing time of the system, which has its own temperature dependence that can be difficult to predict *a priori*, just as in classical Monte Carlo. Indeed, the way glassy or other systems become difficult to simulate is by a divergence of this time at low temperature. But for a conventionally thermalizing system with a mixing time that diverges at worst as a power-law in β, the CKG approach is expected to achieve rapid convergence to the Gibbs ensemble. Its efficiency is related to the quasi-locality of the Lindblad operators, which are concentrated within a range that scales with β.

An insight enabling the result is that one can design dynamics satisfying a detailed balance condition without requiring accurate estimates of energies, which are expensive; this is one reason the CKG approach becomes more efficient for large systems than prior efforts based on simple physical principles, such as that of Davies [7]. There remain a number of steps toward practical implementation, some of which have been studied in subsequent work (for an accessible recent review, see [8]). Fortunately, some of these steps involve interesting physics questions. The precise convergence rate for various systems provides an operational definition of mixing time that may be somewhat different than the mixing time under the original dynamics of the system, similar to how, in classical Monte Carlo methods, one frequently chooses an artificial dynamical evolution that satisfies detailed balance but mixes more rapidly than the physical evolution. Understanding the efficiency of CKG-type algorithms in various states of matter is a major need. But there is considerable hope, among the growing community of researchers seeking to apply quantum computers to challenging materials problems, that efficient simulation of low-temperature quantum matter will arrive sooner rather than later.

The author acknowledges programs of the Simons Institute for the Theory of Computing and the Challenge Institute for Quantum Computation for enabling discussions of the material in this commentary.

References

[1] F. Barahona, J. Phys. A: Math. Gen. **15**, 3241 (1982).

[2] T. Kibble, Physics Reports **67**, 183 (1980).

[3] W. Zurek, Physics Reports **276**, 177 (1996).

[4] K. J. Satzinger, Y.-J. Liu, A. Smith, C. Knapp, *et al.*, Science **374**, 1237–1241 (2021).

[5] E. Rosenberg, T. Andersen, R. Samajdar, A. Petukhov, *et al.*, Science **384**, 48 (2024).

[6] C.-F. Chen, M. J. Kastoryano, and A. Gilyén, "An efficient and exact noncommutative quantum gibbs sampler," (2023), arXiv:2311.09207 [quant-ph] .

[7] E. B. Davies, Communications in Mathematical Physics **39**, 91 (1974).

[8] L. Lin, "Dissipative preparation of many-body quantum states: Towards practical quantum advantage," (2025), arXiv:2505.21308 [quant-ph] .